Slim Shading in Ad Auctions:

Adjustment of Bidding Strategies to First-Price Rules

Miguel Alcobendas
Yahoo Research

Robert Zeithammer
UCLA

February 14, 2023

Abstract: We document the response of bidders to a switch in auction pricing rules by a platform for selling online advertising impressions. The platform switched from a second-price auction to a first-price auction, so the same bidder bidding to show the same creative in the same location on the same webpage should bid less after the switch than before the switch. We show that bids indeed decline after the switch, but they do not decline enough given the actual competition each bidder is facing after the switch. To measure whether bids declined enough, we propose a nonparametric estimator of a lower bound on the bidder's valuation underlying each post-switch bid. Bids did not decline enough in that the estimated bounds substantially exceed the pre-switch valuations of showing the same creative. We find evidence of an incomplete and slow downward adjustment in bid magnitude over the period of months, whereby bids remain insufficiently shaded for about three quarters of the creatives we analyze even three months after the switch. Our results have important implications for analysis of bidding in first-price auctions, interpreting the revenue effects of changes in auction rules, and analysis of short-run A/B tests of different pricing rules.

Contact: Robert Zeithammer, rzeitham@ucla.edu

1. Introduction

Most platforms for real-time bidding on displaying online advertising have recently shifted from the second-price to the first-price auction format (Despotakis et al 2019, Choi et al 2020). Many of these transitions happened en bloc, meaning that the entire platform switched to the new pricing rule at one time after pre-announcing the switch date broadly. We analyze one such transition from early 2019. The platform that provided us with data prefers to remain anonymous, so we cannot disclose the exact date and instead measure time relative to the date on which the switch occurred throughout the paper. Our main question is how well and how fast did the bidders adjust their strategies to the new auction format. We conduct two analyses of the market's response to this switch: a "macro" analysis of the effect of the switch on the publisher's revenue and average price of displaying an ad, and a "micro" analysis that zooms in on individual-bidder bidding strategies associated with specific long-running creatives that span the time of the switch.

Our first "macro" analysis uses the approach of Gobillon and Magnac (2016) to compare the post-switch revenues and transaction prices of a large website served by the platform that switched pricing rules to counter-factual revenues predicted based on several comparable websites owned by the same publisher but served by a platform that did not switch rules. We find that both publisher revenue and ad prices increased temporarily, returning to their pre-switch levels about five weeks after the switch. This pattern is in line with results reported by Goke et al. (2022), who show analogously transitory increases in the prices of advertising using data from different online publishers served by a different platform which also switched auction formats in 2019.

Our second "micro" analysis asks whether the eventual reversion to pre-switch levels found by the first analysis can be attributed to a sufficient downward adjustment of bidding strategies in accordance with auction theory and the Goke et al. (2022) interpretation of their macro findings in an analogous switch of auction rules (Goke et al. did not conduct auction-level analysis, but they propose their macro revenue results may result from sufficient adjustment of bidding strategies – a potential explanation we can test directly). We now briefly outline what standard theory predicts about the situation.

First, bids should decline after the switch, ceteris paribus, because auction winners suddenly have to pay their own bid instead of the second highest bid. The bid-reduction below valuation (which is the dominant-strategy bid under second-price rules) is called "bid shading" in both industry (e.g. Sluis 2019) and academia (e.g. Battigalli and Siniscalchi 2003). Second, previous analyses of bidders' adjustment to new rules (e.g. Doraszelski, Lewis and Pakes 2018) predict that different bidders adjust at different speeds. The market may thus take some time to converge to the new equilibrium, with the transitional period characterized by higher prices due to insufficient adjustment by some bidders. Note that the adjustment is not simple because bidding in first-price auctions places a higher computational and informational burden on bidders than bidding under second-price rules in that the bidders now need to take the intensity of competition into account when shading their bids. Finally, the individual bidding strategies should converge to each bidder playing their best response to the competition they face so that the collection of bidding strategies constitutes an equilibrium. In a first price auction, this means each bidder should shade their bid just enough to balance the chance of winning against the surplus gained upon winning.

We ask three questions motivated by the above theoretical predictions: Did the bids on comparable opportunities fall? Specifically, did the bids by the same bidder bidding to show the same creative in the same location on the same webpage fall? If yes, did the bids fall enough given the actual competition at the time? And if the bids fell enough, how long did the transition take? To answer these questions, our second "micro" analysis relies on auction-level data about several creatives that generated a consistently high volume of bidding throughout our 4-month data period.

We find that bids indeed declined after the switch for most of the long-running creatives we study, with an average decline of about 5 percent. But was this decline sufficient from the perspective of the individual bidders at the time? To answer this key question about the magnitude of observed bid-shading, we propose a simple nonparametric estimator of the lower bound on advertiser valuations in first-price sealed-bid auctions. The estimator is based on two key assumptions: 1) each bidder prefers the bid he submits to counter-factual lower bids, and 2) each bidder conditions his belief about momentary competition only on common observables and his own valuation of the opportunity. We show how these two assumptions can be used to construct a non-parametric lower bound on the valuation behind every bid in our data without assuming of bidder symmetry, valuation independence, or equilibrium bidding – convenient but highly unrealistic assumptions in our setting. To assess the estimator's precision, we adopt a Bayesian perspective and derive closed-form standard errors based on the distribution of a ratio of two Beta distributions.

We use our bounds estimates to answer our second question ("did the bids fall enough given the actual competition at the time?"), as follows: Before the switch, the truth-revealing property of the second price auction allows us to directly equate valuations with bids. Comparing the post-switch valuation lower bounds to pre-switch valuation magnitudes allows us to conservatively detect insufficient adjustment whenever the former exceeds the latter. And we indeed find that the bid shading was insufficient for a vast majority of bidders and creatives: observed bids on the average creative imply that the bidder representing that creative bid *as if* the switch from second-price to first-price rules increased his valuation of an impression at least 30 percent.

Our proposed measure provides information about bidder's bidding strategy above and beyond the prices the bidder pays for impressions (as measured by CPM - cost per thousand). While prices increased for all but one creative after the switch, such an increase can be attributed not only to insufficient shading by the focal bidder, but also to a violation of revenue equivalence or to an intensification of competition after the switch in the form of additional entry or insufficient shading by competing bidders. Our bounds estimator is designed to control for these alternative explanations of the higher CPM by controlling for the actual competition within each post-switch auction, and by analyzing each bidder's perspective in each auction separately.

Not only do we find insufficient shading right after the switch occurred, but we are also unable to detect any downward trend in the lower bound on valuations expressed as a percentage of pre-switch valuation. In summary, we conclude that the bidders we study took more than three months to adjust to the new pricing rule if they ever adjusted at all. Every one of the three multi-creative bidders in our data has at least one creative with consistently insufficient shading throughout the data period, so we also do not find evidence of heterogeneity in bidding sophistication across bidders.

We do not pinpoint the reason why the bidders we study do not shade their bids on their long-running creatives enough, but we do rule several possible explanations out. First, bids are not systematically increasing after the switch, so the seemingly higher valuations revealed by our estimator are not due to trend seasonality in valuations. Second, the standard deviation of bids does not increase over time, so our data does not support explanations built around creative-specific valuations gradually drifting apart akin to a random walk. Third, volume of impressions purchased increases after the switch for most of the creatives we study, so the seemingly higher valuations revealed by our estimator cannot be explained by a rational forward-looking bidding strategy adjusting to the lower opportunity costs of winning due to higher prices. Fourth, it is unlikely that our robotic bidders are risk-averse like human bidders in laboratory settings who also do not shade their bids in first-price auctions as much as standard theory would predict (e.g. Cox, Smith and Walker 1988, Bajari and Hortaçsu 2005). Finally, our findings are not driven by a particular definition of "pre-switch valuation" as the average bid on the creative during the month before the switch: if we conservatively take the highest pre-switch average weekly bid of the four pre-switch weeks available as the estimate of pre-switch valuation of the impression, we still find evidence of insufficient bid-shading among most of the creatives we study. One explanation we cannot rule out is that the bidders we study are simply not paying enough attention to their bidding strategy on their long-running creatives.

Regarding the implications of our results for the broader market for digital display advertising, it is important to note that our sample selection to long-running creatives with lots of bids throughout the period makes our conclusions limited to the bidders and creatives we study, and should not be taken as a general characterization of the entire market. Instead, we view our main contributions to be the development of a simple nonparametric estimator of the lower bound on advertiser valuations in first-price sealed-bid auctions, and using the estimator to expose of the difficulty of bidding-strategy adjustment in the early days of selling online advertising by first-price auctions. The slow (if any) adjustment we document suggests that the transitory increase in revenues found in our "macro" analysis of the same publisher cannot be explained by bidders adjusting their bidding strategies in accordance to theory on the same time scale. Instead, it seems other adjustments must have also taken place, for example an adjustment of bidding strategies on only new creatives by the bidders we study, or an adjustment of bidding and/or entry strategies by some other bidders. Finally, the slow (if any) adjustment also raises the challenge of interpreting short-run A/B tests of large changes to an auction marketplace, such as changes in the pricing rule.

2. Literature review

The programmatic digital advertising market represents a large and growing proportion of all advertising spending in the U.S., so understanding its inner working is central to Marketing. Our main contribution is empirical in that we contribute to a better understanding of how this important market operates. Specifically, we contribute to the growing literature on real-time bidding (RTB) on display advertising (see Choi et al 2020 for a recent review), which is the dominant form of digital advertising today, having surpassed search advertising in dollar terms in 2016. Most papers in the literature keep the market rules fixed and analyze the RTB market from the auctioneer's perspective focusing on, for example, the problem of setting reserve prices (Choi and Mela 2018), the decision whether to use soft floors (Zeithammer 2019), or the strategy for incorporating dynamic ad sequencing (Rafieian 2019). An exception is Despotakis, Ravi and Sayedi (2021) who claim that the switch we are studying happened as a result of header bidding. In contrast to Despotakis, Ravi and Sayedi (2021), our paper is empirical and we do not attempt to model why the switch occurred, but merely measure the adjustment of bidding strategies to it.

The closest existing work is the working paper by Goke et al (2022), who consider daily prices of display advertising aggregated to the publisher-bidder level, and analyze the quasi experiment arising from the fact that not all publishers using one particular exchange switched to the first-price rules at the same time. They find a multi-month increase in prices consistent with insufficient bid shading, and a general reversion of prices back down to pre-switch levels after about 60 days. Our first ("macro") analysis conceptually replicates these findings in our setting using a linear factor model approach. In contrast to Goke et al (2022), our second analysis takes a more micro approach, and focuses on analyzing how individual bidders bid in individual auctions. This second analysis allows us to test whether the transitory nature of the price and revenue increases after the switch is due to an adjustment of bidding strategies on the same time scale. Surprisingly, we do not find sufficient adjustment of the bidding strategies, so we propose other types of adjustments must be responsible for the reversion of prices and revenues back to pre-switch levels.

Our main methodological contribution is the non-parametric estimator of the lower bound based on a hyper-local estimate of competition. The estimator builds on the seminal work of Guerre, Perigne, and Vuong (2000), but does not construct a point estimate of the valuation as is customary in the econometrics literature about first-price auctions (e.g. Athey and Haile 2002). Instead, only a bound is inferred as in Hortaçsu and McAdams (2010) who used inequalities analogous to our first main assumption to construct an estimate of bid-shading in Turkish treasury auctions. The theoretical advantage of our weaker assumption over all previous papers is that we do not need to assume global optimization by every bidder. The computational advantage of our weaker assumption is that our bound estimator needs only an estimate of the cumulative distribution function of competing bids instead of also requiring an estimate of the probability density that would be needed for a point estimate, which is notoriously more difficult to estimate precisely.

While weaker than the global optimality assumption, our assumption about bidder rationality is stronger than that in Haile and Tamer (2003), who analyze an English auction and only assume that bidders 1) "neither bid more than their valuations" and 2) do not "let an opponent win at a price they would be willing to beat". Because we study a sealed-bid auction, their second assumption is not useful in our context. Moreover, our first assumption is stronger than theirs in requiring partial optimization, and also implying that we need our bidders to have probability beliefs about their chances of winning at different bid-levels. We specify the needed granularity of these beliefs in one of our main assumptions.

Unlike Hortaçsu and McAdams (2010) who use analogous inequalities to analyze treasury auctions and Chan and Park (2015) who use a conceptually similar "locally envy-free" inequality to analyze search-auction bidding data, our approach does not require any assumption about equilibrium bidding. Instead, we consider the bidding problem separately from the perspective of each bidder we analyze, taking the competing bidders as given. This is an advantage of our approach for studying a market in transition, such as ours.

Finally, our derivation of standard errors of our estimator cleanly exploits the mathematical properties of the estimator instead of resorting to computationally prohibitive brute-force approaches such as bootstrapping. Assessing the precision of a ratio of two beta distributions has been of recent interest in statistics (Pham-Gia 2000) and already found applications in epidemiology (e.g. Bekker-Nielsen et al 2019). We show how these results can be useful for measuring the precision of our class of bound estimators.

3. Analysis of overall publisher revenue and ad prices

How did the switch in pricing rule impact publisher weekly revenue and prices of advestising exposures? We focus on one large U.S. internet property served exclusively by the platform we study, and use the linear factor model approach to estimate counterfactual values of both variables (revenues and prices). Specifically, we model each variable at the daily level for ten weeks before the switch as a function of the same variable at several other "control" properties that did not switch. The control properties were also served exclusively by the platform our data is from, but they operated in foreign markets that switched to first-price rules later. Apart from property fixed effects, our model also controls for other covariates like number of ad requests the publisher issued each day, and (only in the case of modeling revenue) also the number of impressions sold[1]. We then follow the standard linear factor model approach by Gobillon and Magnac (2016) to calculate the counterfactual variables (revenues and prices) the focal publisher would receive after the switch if its advertising pricing continued to rely on second-price rules.[2] The approach we use is standard, so we relegate the details to a Web Appendix.

[1] Since price = revenue / # impressions, the price model incorporates this covariate on the right-hand side.

[2] The Gobillon and Magnac (2016) method is based on previous work on interactive fixed effects and consists in an iterative procedure which incorporates time latent factors and fixed effects for the treated and untreated sites. We obtained similar outcomes using the matrix completion model described in Athey et al. (2017) and the synthetic control method proposed by Xu (2017), increasing robustness in our results

Figure 1a: Effect of the switch on publisher revenue (treatment effect on the treated)

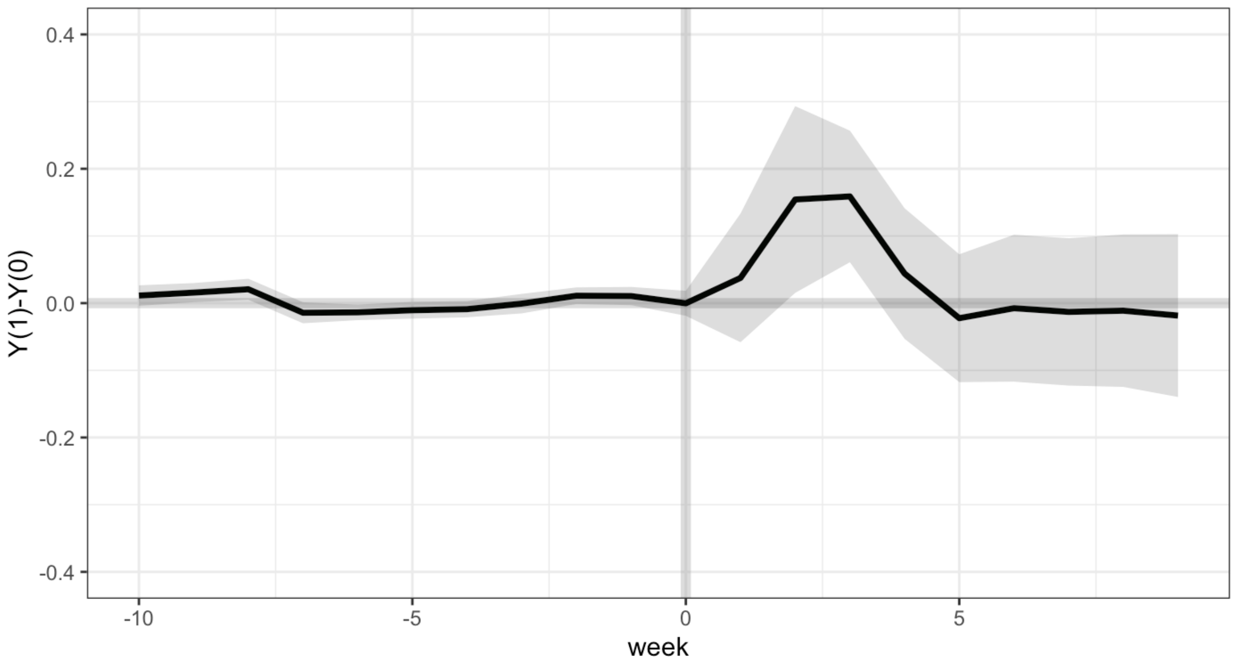

Note to Figure: The dark gray area around the line represents the 95% confidence interval. Figure 1b displays the impact of the transition on the transaction prices (CPM) generated by the focal publisher's website. It is analogous to figure 1a, with only the dependent variable switched.

Figure 1a displays the impact of the transition on the revenue generated by the focal publisher's website. It shows the difference between the observed log revenue (denoted Y(1)), and the log revenue we would expect to get if we continued using a second price mechanism (denoted Y(0)). The black line corresponds to the estimate of the difference between the log revenues and the dark gray area around the line represents the 95% significance level, which is computed using bootstrapping. The x-axis indicates the corresponding week before and after the switch. As expected, the revenue difference between the model and the observed data is not significant before the switch – a proof that our linear factor model is working well. The revenue difference diverges from zero during the first 4 weeks after the switch, although the difference is only significant in weeks 2 and

3. Beyond week 5, the counterfactual suggests that the observed revenue under the first price auction rule is not different from the revenue we would expect to get if we continued using a second price mechanism.

Figure 1b: Effect of the switch on prices of advertising (treatment effect on the treated)

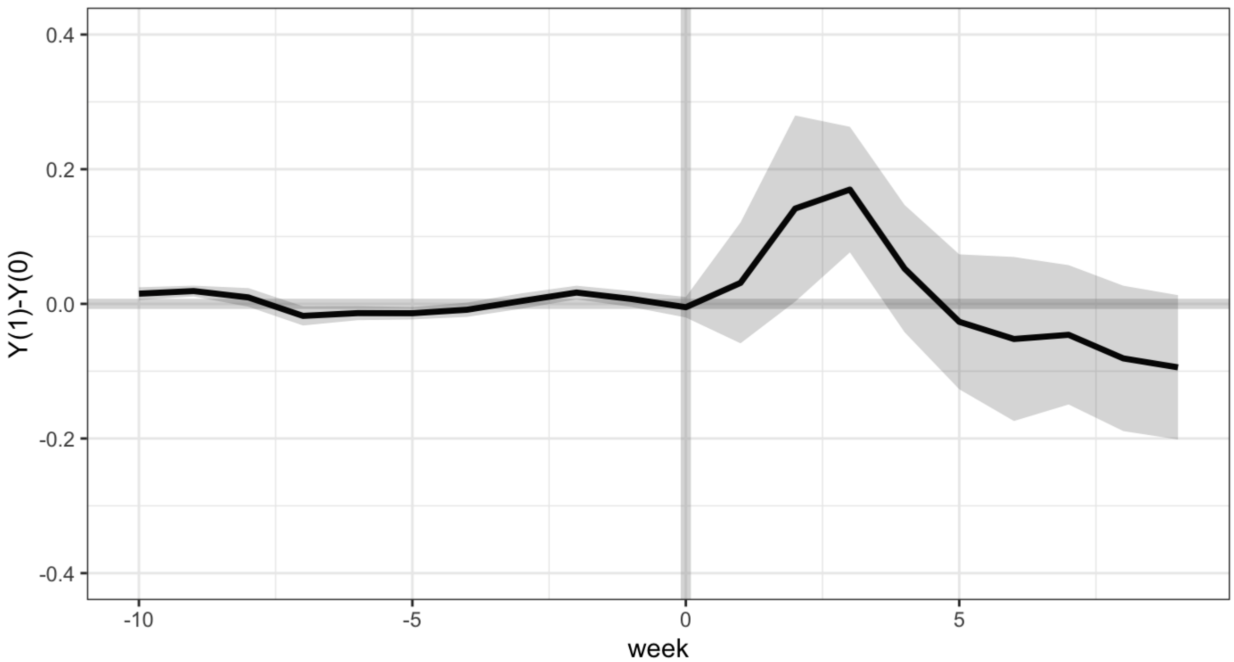

Note to Figure: The dark gray area around the line represents the 95% confidence interval.

The pattern in Figure 1b in line with results reported by Goke et al. (2022), who show a similar transitory increase in prices using data from different online publishers that also experienced a switch in the pricing rule in 2019. They conclude that the average price of advertising returned to pre-switch levels within 60 days of the switch. Figure 1b also suggests that the transitory increase in weekly revenue documented in Figure 1a was driven by higher prices, and not by a higher volume of transactions.

One way to interpret Figures 1a and 1b and the analogous results reported in Goke et al. (2022) is as evidence of bidder first bidding too much, but then learning how to bid under the new rules within a few weeks. In other words, the transitory nature of the post-switch nature of the revenue and price increases can be taken as evidence that the bidders "gradually learned to shade their bids to a level sustained by a rational strategy" (Goke et al. 2022, p. 16). In the rest of this paper, we test this interpretation by zooming in on a few long-running creatives running on the same focal publisher throughout the time we study.

4. Analysis of bidding strategies: data on long-running creatives

The previous section is consistent with the bidders gradually adjusting their bidding strategies to the new pricing rule within a few weeks of the switch. In the auctions we study, the advertisers do not bid directly, but rather via "Demand Side Platforms" (DSPs for short). DSPs pick and choose among the millions of ad opportunities available on the internet on behalf of their advertiser clients, and formulate bids to satisfy the advertisers' campaign goals. Did these institutional bidders actually bid less after the switch as rational theory would predict? If yes, did their bids fall enough? And if the bids fell enough, did the transition take about five weeks as Figure 1 would suggest? To answer these questions as cleanly as possible, we need to hold as many things as possible constant throughout the data period. So we set out to analyze how the same bidder bid on showing the same creative in the same location on the same webpage before and after the switch.

To find such long-running creatives placed over and over by the same bidder, we collected a large random sample of bids on the platform for the most popular ad size and location on the same publisher's website for four months from one month before the switch to three months after the switch. To ensure enough data about each bidder, we focus on the top bidders in terms of bid volume[3], and select their creatives that received enough serious bids throughout the time period. Specifically, we look for creatives that have (within our sample) at least one thousand bids every week during the 4-month data period, with every week's median bid above the reserve price. Fortunately for our goal of analyzing the effect of the switch while holding as many observables as possible constant, the reserve price on the exchange we study stayed around the same level throughout the observation period and did not respond to the switch in any discernible way.

The data selection procedure described above yields 19 million bids on 11 creatives by 4 different bidders – all major players in the industry.[4] We cannot reveal the identities of the advertisers, only that the creatives we study advertise a range of consumer products and services, including consumer durables, home-goods retailers, and financial services. To protect the identity of the bidders, we label them with letters of the alphabet in no particular order. Each bidder's creative is then assigned a number in no particular order. Thus, our unit of observation is a particular bidder (e.g. C) bidding on a particular creative (e.g. 2), labeled "C2". For brevity, we occasionally refer to the bidder-creative pair simply as creative, e.g. "creative C2". Table 1 summarizes the bidding data in our sample by month along with the average cost per thousand impressions (CPM) defined as the amount of money paid divided by the number of thousands of auctions won.

---

[3] One other large bidder is associated with the platform at the corporate level. We exclude this bidder from our analysis because their incentives may be different from the simple surplus maximization we assume in our model.

[4] Our sample is ideal for measuring the adjustment of bidding strategies over time, but it is in no way representative of the typical creative in the market – most last less time, and there are also many smaller bidders who bid too sporadically to be useful for our analysis. Our sample selection thus makes our conclusions limited to the bidders and creatives we study, and should not be taken as a general characterization of the entire market.

Table 1: Data summary

Creative	Number of bids	average bid by month (% of month before switch)			average CPM by month (% of month before switch)		
		1 after	2 after	3 after	1 after	2 after	3 after
A1	968,238	106%	102%	113%	157%	149%	164%
B1	778,412	103%	107%	94%	120%	124%	105%
B2	6,612,124	113%	111%	89%	135%	138%	110%
B3	505,714	102%	94%	76%	124%	119%	101%
B4	123,187	85%	90%	96%	205%	203%	205%
C1	4,253,188	90%	89%	94%	139%	136%	137%
C2	3,254,961	84%	85%	86%	139%	137%	137%
D1	214,450	71%	69%	67%	94%	95%	92%
D2	159,370	124%	104%	101%	165%	137%	130%
D3	1,612,881	57%	85%	97%	86%	126%	142%
D4	487,630	100%	125%	133%	114%	142%	151%
Average	**1,724,560**	**94%**	**96%**	**95%**	**134%**	**137%**	**134%**

To motivate the development of our bound estimator, we now briefly interpret the data summary in Table 1, starting with bid amounts. The average decrease of 5 percent conceals a lot of variation: a month after the switch, bids decreased substantially (about 20 percent) for about two thirds of the creatives we analyze, and increased for the remaining creatives mostly to a lesser amount. We conclude that, consistent with theory, the bidders we study did generally shade their bids down in response to the switch in pricing rule. However, we can also already rule out a pure bidder heterogeneity story as a potential explanation of differences in adjustment to the price-rule switch: the creatives of bidder D include both extremes of the post-switch difference (highest increase for D2, highest decrease for D1). Analogously, bidder B reduced bids on one creative (B4) while increasing them on another one (B2) and leaving bids essentially unchanged on the other two.

Figure 2: Bids on three selected creatives (4-hour moving average)

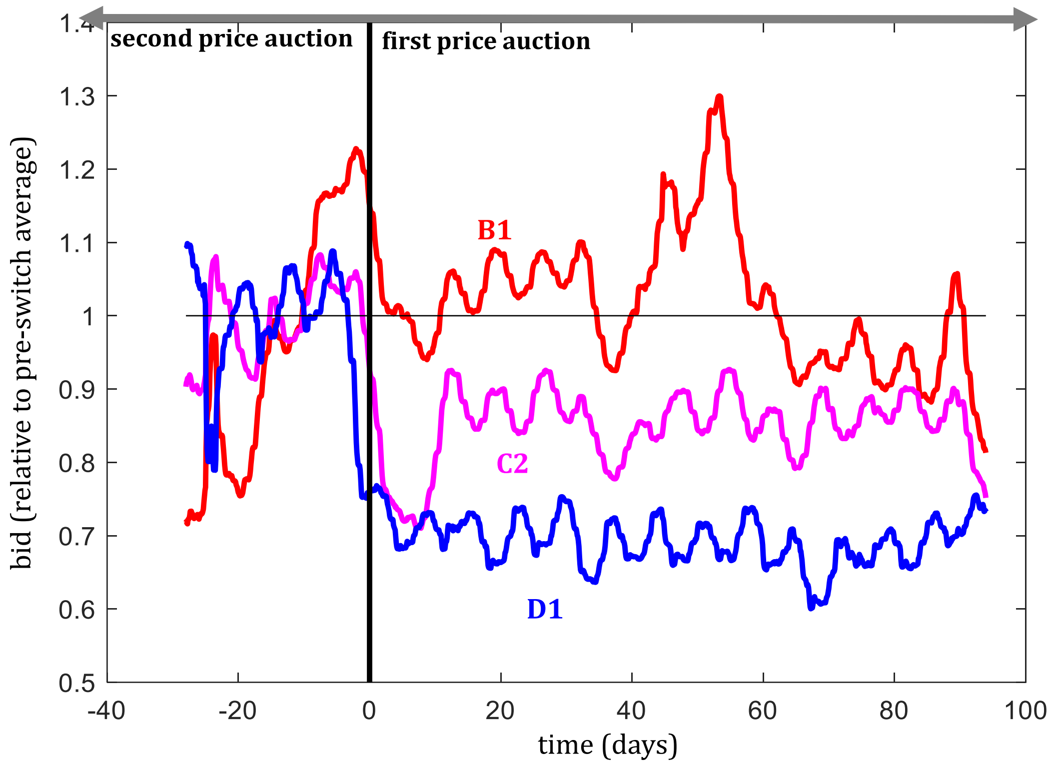

Figure 2 plots the 4-hour moving average of bids on three different creatives by three different bidders (see color coding in Table 1), and it shows that C2 and D1 definitely adjusted to a fairly constant lower level after the switch while B1 did not do so. However, it also seems that B1 did eventually adjust downward as well, albeit only after two months. An interesting pattern also visible from Figure 2 its that the weekly fluctuation in bid amount seems synchronized between B1 and C2, but runs opposite for D1. In other words, B1 and C2 bid more on weekends while D1 bids less.

While it is clear from Table 1 and Figure 2 that several creatives shaded bids in the same direction as theory would predict, a question remains whether they shaded their bids enough given their valuations of the impressions and the competition they faced. If one believes that revenue equivalence between first- and second-price auctions applies at least approximately in this setting, then the cost per thousand impressions (CPM) data in Table 1 suggests that the adjustment was mostly insufficient: the effective price of an impression increased substantially and permanently for all but one creative (D1). However, there are at least two possible explanations of the CPM increase among our long-running creatives alternative to insufficient shading by our bidders: First, first-price auctions may revenue dominate second-price auctions in our setting for several reasons.[5] Second, even if the bidders we study shaded sufficiently given the situation they are in, competition may have intensified due to additional entry or insufficient shading by competing bidders. Our bounds estimator is designed to disentangle the insufficient shading explanation from these two alternative explanations of the higher CPM by both controlling for actual local competition and by directly analyzing the bidding problem from each bidder's perspective instead of only looking at revenue.

5. Theory: Deriving a lower bound on the bidder's valuation from an observed bid in a first-price sealed-bid auction

Consider a single first-price sealed-bid auction (FPSB) with a reserve price of R, and let b be a bid submitted in the auction by one participating bidder. The bidder believes his probability of winning the auction is captured by a cumulative distribution function $G(b)$ on $[R, \infty)$, and values winning the auction some amount *x.* Note that if this auction is actually one of many auctions selling similar things (e.g. opportunities to show the same advertisement to a similar target audience), the valuation *x* reflects both the option value of

[5] Bidder symmetry in valuations required for revenue equivalence is unlikely to hold in our setting because different bidders have different quality and quantity of information about each opportunity. In addition, as long as the DSPs hold internal auctions to decide which advertiser to advance to the auction we study, then the results of Despotakis et al (2019) would suggest that the first-price auction is more profitable.

other opportunities to win the same object and the number of identical objects the bidder wants to win. We thus abstract from dynamic auction issues, assuming the bidder has solved the relevant dynamic program to compute his *net* valuation of a single opportunity when facing a stream of opportunities. See the Appendix for a formalized argument why a *net* valuation summarizes the dynamic future opportunities adequately in this context, as well as for additional references. We assume the following one-directional rationality:

Assumption* 1 *(Observed bid preferred to lower bids): The bidder prefers the bid he submits b to counter-factual lower bids involving additional shading $s > 0$:

$$G(b)(x-b) > G(b-s)(x-b+s) \tag{1}$$

Note that the above Assumption 1 is weaker than the standard assumption in auction econometrics following Guerre, Perigne, and Vuong (2000) that a bidder fully solves the optimization problem $b = \underset{y}{argmax}\, G(y)(x-y)$. Specifically, we do not assume that the analogous inequality for higher bids, namely $G(b)(x-b) > G(b+s)(x-b-s)$, also holds, which contrasts our work with previous approaches that use inequalities to derive bounds on valuations in related settings (e.g. Hortaçsu and McAdams 2010, Chan and Park 2015).[6]

If we knew the focal bidder's beliefs about the G probability function in equation 1, we could simply rearrange the inequality to derive a lower bound on x as a percentage of b as follows:

$$\frac{x}{b} > 1 + \left(\frac{s}{b}\right)\left(\frac{G(b-s)}{G(b)-G(b-s)}\right) \tag{2}$$

[6] Hortaçsu and McAdams (2010) use a similar inequality in their analysis of bidding in Turkish treasury auctions, but they do assume both complete optimization by each bidder, as well as equilibrium and independence in valuations across bidders. Chan and Park (2015) use a conceptually similar "locally envy-free" inequality to analyze search-auction bidding data, but they again require equilibrium for their approach.

However, the bidder's beliefs about G are not part of our data, so we need to make additional assumptions to make the calculation of inequality 2 feasible. Following Guerre, Perigne, and Vuong (2000), most previous papers about the econometrics of a first-price sealed-bid auction enable the analyst's inference about G by assuming the bidders are symmetric (after conditioning on observables) and the market is in a symmetric pure-strategy equilibrium (implicitly assuming full optimization by each bidder). Then, all observed bids can be interpreted as interchangeable and pooling across bidders and auctions yields an empirical estimate of G each bidder is facing in each auction.

We cannot assume equilibrium bidding because our goal is to analyze a market in transition, and all bidders may not adjust their strategies in the same way or at the same speed. To accommodate our market in transition, we analyze each bidding problem separately from the position of each bidder, simply taking the observed behavior of competing bidders as given. Besides not requiring any equilibrium to hold, this approach also does not rely on bidder symmetry.

So how do we assume each bidder to formulate his beliefs about G? The bidder obviously conditions on various observables known to both him and the analyst. In our digital advertising context, such observables include the ad to be shown, the internet property being auctioned, and the moment in time in which the auction occurs. It is obvious that we, as the analysts, also need to condition our measurement of G on these common observables. Regarding variables observed by the bidder but not observed by the analyst, such as the information in the bidder's cookie data regarding the quality of the particular exposure opportunity, we assume that the valuation x is a sufficient statistic:

Assumption* 2 *(Bidder beliefs about competition): The bidder's private valuation x is a sufficient statistic about the impact of variables observed by the bidder but unobserved by the analyst on the bidder's beliefs about his probability of winning the auction.

In other words, if there are two auctions indexed *i*=1,2 with the same observables Z but different unobservables $U_1 \neq U_2$ such that $x(Z, U_1) = x(Z, U_2)$, then we assume that for every possible bid-level $y < x$, the bidder believes that his chances of winning are the same: $E[Pr(win|y, Z, U_1)] = E[Pr(win|y, Z, U_2)]$. This assumption is not innocuous: suppose the unobservable variable $1 = U_1 \neq U_2 = 0$ is an indicator of a crazy high-bid competitor participating in the auction. Since such a competitor is known to just bid high no matter what, the competitor's participation may not change the beliefs of our focal bidder regarding his valuation of the potential ad exposure, so $x(Z, U_1) = x(Z, U_2)$ may hold. However, the focal competitor would certainly believe that $E[Pr(win|y, Z, 1)] < E[Pr(win|y, Z, 0)]$. Assumption 2 assumes that the bidders we study do not have access to auction-level unobserved signals of this kind. In other words, we assume the bidder's unobserved information pertains mostly to how much they value the exposure, not to the intensity of competition they are about to face in the individual auction.

If we knew the bidder's valuation *x* in each auction, Assumption 2 would allow us to pool across auctions in which the bidders has the same *x.* Learning about *x* is the primary objective of the whole exercise, so we cannot condition on it directly. However, one more assumption will allow us to exploit the one-to-one relationship between bids and valuations:

Assumption* 3 *(Bids increasing in valuations): Conditional on observables *Z,* the bidder's bidding strategy $\beta(x|Z)$ is strictly increasing in his valuation *x.*

In other papers on the econometrics of first-price sealed-bid auction, Assumption 3 usually tacitly follows from the assumption of symmetric equilibrium and some sort of regularity of the joint distribution of valuations or signals. Here, we simply assume that as the valuation of an impression increases, the competition our bidders face does not become so much weaker that their best-response bids to it would actually fall. Given assumption 3, conditioning the belief G on the observed bid-level is the same as conditioning it on the unobserved valuation that gave rise to the bid. We will denote this expected probability of winning as $G(bid|b) \equiv E[Pr(win|bid, Z, U_i)]$ where $b = \beta(x(Z, U_i)|Z)$. Note that this notation suppresses the dependence of G on observables Z to avoid cluttering our notation. Li, Perrigne and Vuong (2002) show that this type of conditioning of beliefs on valuations arises naturally in an equilibrium of a symmetric first-price sealed-bid auction under affiliated private values. Our assumptions 2 & 3 together thus accommodate but do not assume affiliated private values. Importantly, the two assumptions allow valuations of different bidders for the same opportunity to be interdependent, and we will show below that a positive correlation of bids across bidders is an important feature of our data.

Assumptions 2 and 3 allow us to substitute $G(.|b)$ for G in inequality 2, and define the lower bound on valuation $LBV(b)$ as:

$$x > LBV(b) \equiv b + \max_{s\in[0,b-R]} s\left(\frac{G(b-s|b)}{G(b|b) - G(b-s|b)}\right) \tag{3}$$

which follows from the fact that the bidder prefers the observed bid to all counterfactual bids on $[R, b)$, i.e. all feasible s that result in a positive probability of winning. Nekipelov, Syrgkanis, and Tardos (2015) show that under mild assumptions, an analogous bound can apply even when the bidder is merely using no-regret learning as opposed to responding to a known G as we posit. We now turn to our empirical strategy for estimating $LBV(b)$ with our data, and measuring the precision of our estimate.

6. Non-parametric estimation of the lower bound on valuations

Fix a bidder and an auction the bidder participates in with a bid b_0 and denote the highest competing bid Y_0. Suppose the analyst observes K additional auctions indexed by k=1,2,…, K with identical observables in which the focal bidder also submitted a bid. For each of these auctions, the analyst observes the focal bidder's bid b_k and the highest competing bid Y_k. Given the shared identical observables, the K+1 auctions thus differ only in the private valuation component of the bidder's valuation. Therefore, the distribution of Y_k in the subset of the K+1 auctions such that $b_k = b_0$ is the precisely the distribution $G(.|b_0)$ we need to calculate $LBV(b)$ according to equation 3. If, for every focal auction, we had a large set of auctions for which the observables were truly identical and $b_k = b_0$, we could simply estimate $G(.|b_0)$ as the empirical cumulative distribution function of the corresponding Y_k from this set. In practice, we need to resort to observables that are *approximately* identical and bids b_k that are *approximately* equal to b_0, as we explain next. Some sort of continuity of beliefs in valuations and observables is clearly required here, but we do not think additional definitional formalism is necessary to get our idea across.

In our application to display ad auctions, the focal bidder is a particular DSP (Demand Side Platform, see Data section for institutional details) bidding to show one particular creative on a particular property at a particular moment in time. We consider other auctions attended by the same bidder to have approximately identical observables when those auctions involve the bidder bidding to show the same creative on the same property within a few hours of the focal auction. From this set of auctions with approximately identical observables, we then select the auctions with a bid amount b_k within a few cents of b_0, and use the empirical cumulative distribution function of the associated competing bids Y_k to estimate the $G(.|b_0)$ in calculating the $LBV(b)$ according to equation 3.

Figure 3: Illustration of how we define local competition

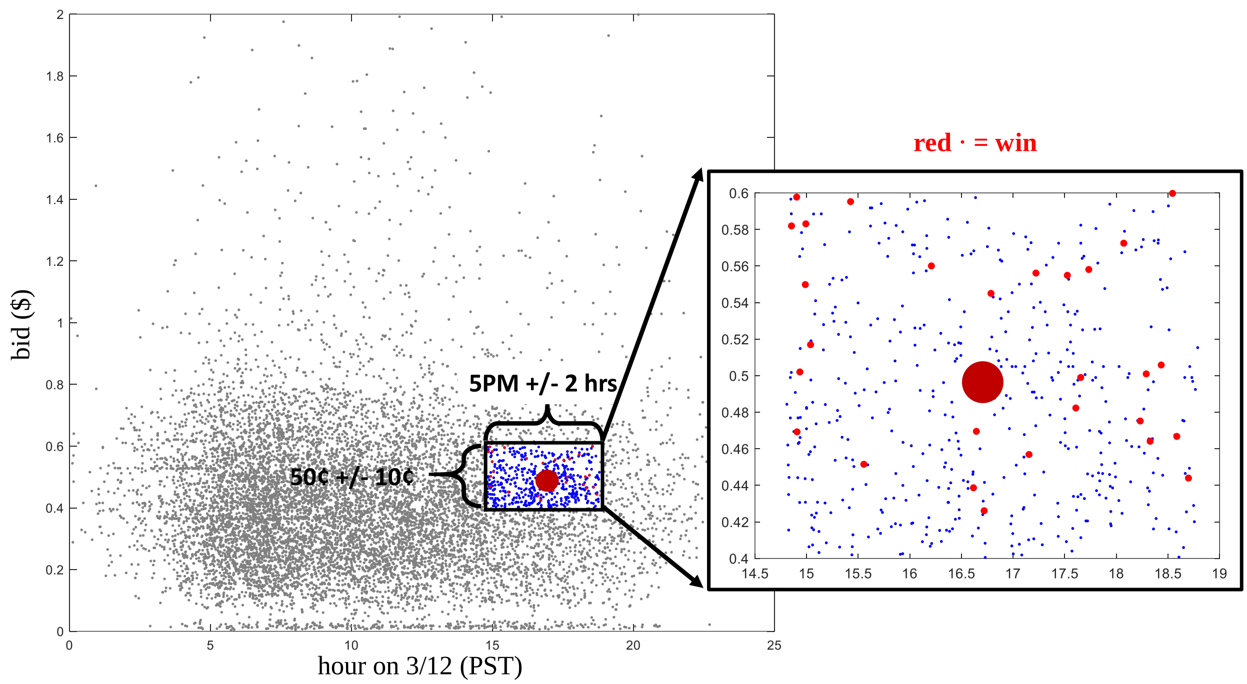

Figure 3 illustrates our approach on a particular focal auction held at 5 PM on 3/12/2019, in which the bidder offered b_0 =50 cents (all bids are on the CPM basis, i.e. per thousand impressions). That day, that bidder made 10,757 bids to show the same creative as in the focal auction on the same property as the focal property, shown as grey dots in Figure 3. To estimate the competition faced in the focal auction, we focus on only the 476 auctions "near" the focal auction in the sense of occurring within two hours and in the sense of the focal bidder bidding within 10 cents of the focal amount. Figure 3 zooms in on these 476 "nearby" auctions, and shows in red the 28 bids that exceeded their respective highest competing bid, and thus resulted in a win. However, we do not pay any attention to the winning or losing status of each bid by our focal bidder, i.e. whether or not $b_k > Y_k$. Instead, we estimate the cumulative distribution function of the 476 *competing* bids Y_k in the nearby auctions to characterize the $G(.|b_0)$.

By conditioning the relevant competition on the focal bid-level, this approach is most similar to Li, Perrigne and Vuong (2002), who studied the equilibrium bids in a first-price sealed-bid auction under symmetric affiliated private values. In contrast, we study only the best response of one bidder to the competition he is facing at the moment, so we have to re-compute the G separately from each bidder's perspective for each auction. Thanks to this local focus, we can accommodate asymmetry in valuation distributions and we do not require any equilibrium to hold. Assumptions 2 and 3 are clearly sufficient for our approach to be valid, and they may even be necessary.

Note that our approach requires data on the highest *competing* bids Y_k that share the same observables (creative, property and time) and unobservables (bid level), not only a set of bids b_k that share the same observables and their winning status. Suppose a bidder conditions on observables and constructs the cumulative probability of winning at different bid-levels $\tilde{G}(b)$. In Figure 3, such an approach would simply look at the proportion of red dots as a function of bid-level b_k in a neighborhood of 5 PM. Such "naïve" $\tilde{G}(b)$ would obviously be very close to our $G(b|b_0)$ as long as $b \approx b_0$. However, as long as the private valuations are affiliated and b_k thus positively correlated with Y_k, $\tilde{G}(b)$ would be greater than $G(b|b_0)$ for $b \ll b_0$. In words, actual lower bids by the focal bidder tend to win more often than would be the chance of winning with a counterfactually lower bid because lower actual bids are associated with weaker competition. The difference between $\tilde{G}(b)$ and $G(b|b_0)$ is not only a theoretical possibility: Figure A1 in the Appendix shows both the naïve and the correct probability curves at two different own-bid levels on a particular day for one of our creatives. One take-away from Figure A1 is that value-affiliation is a significant issue in our data, and our method's ability to handle affiliation is thus a key strength.

While the distinction between $\tilde{G}(b)$ and $G(b|b_0)$ outlined in this paragraph poses no problem for us (we do have data on Y_k), it highlights the problem bidders face in learning their own counterfactual probabilities of winning, and thus refining their strategies. Instead of merely logging their wins while following a given strategy, they also need to constantly experiment with different bid-levels for the same opportunity.

7. Parametric approach to measuring the precision of LBV estimates

Our estimate of LBV involves known quantities (*b, s,* and R) and the empirical estimate of $G(.|b_0)$, which is clearly estimated with error and appears multiple times in the *LBV* formula. To assess the precision of our *LBV* estimate, we need to evaluate the precision of the key ratio term $\frac{G(b-s|b)}{G(b|b)-G(b-s|b)}$ in equation 3. Notice that this ratio is a ratio of two proportions: the numerator is the proportion p_1 of the local competing bids that are below b-s, and the denominator is the proportion p_2 of the local competing bids that are between b and b-s. To assess the standard error on the ratio of two proportions, we take the Bayesian perspective following Pham-Gia (2000) and adopt as analysts a diffuse Beta prior conjugate to the binomial likelihood of a proportion. Specifically, we assume that $p_i \sim Beta\left(\alpha_{i,0}, \beta_{i,0}\right)$, where $\alpha_{i,0} + \beta_{i,0} = 10$ and the prior mean $\frac{\alpha_{i,0}}{\alpha_{i,0}+\beta_{i,0}}$ is centered on the empirical estimate of the relevant proportion based on a 24 hour period preceding the local competition window used in the estimation of G. In words, we set our prior beliefs separately on each proportion p_i as if we observed 10 binary success/failure observations in the past, and the proportion of successes among these 10 prior "pseudo-observations" was the empirical estimate of the proportion based on recent history. Given this prior belief, it is well known (thanks to conjugacy) that the posterior distribution of each p_i given S_i successes from N_i observations is again from the Beta family. It is less well known but extremely useful for our purposes that the posterior moments of the ratio of p_1/p_2 have the following simple closed form:

$$E[(p_1/p_2)^k] = \frac{\left(\alpha_{1,0} + S_1\right)_k}{\left(\alpha_{1,0} + \beta_{1,0} + N_1\right)_k} \cdot \frac{\left(\alpha_{2,0} + \beta_{2,0} + N_2 - k\right)_k}{\left(\alpha_{2,0} + S_2 - k\right)_k} \qquad (4)$$

where $(x)_k = \frac{\Gamma(x+k)}{\Gamma(x)}$ is the Pochhammer symbol (Pham-Gia 2000). After updating our prior beliefs (specified above) with N observations of competing bids of which S_1 were below b-s and S_2 were between b and b-s, we thus use the following posterior mean as our estimate of the key ratio:

$$E\left(\frac{G(b-s|b)}{G(b|b)-G(b-s|b)}\right) = \left(\frac{9+N}{10+N}\right)\left(\frac{\alpha_{1,0}+S_1}{\alpha_{2,0}+S_2-1}\right) \qquad (5)$$

The posterior mean obviously approaches the raw empirical estimate S_1/S_2 as N increases to infinity. The key for assessing our uncertainty about this estimate is the closed-form second moment, which in turn implies:

$$Var\left(\frac{G(b-s|b)}{G(b|b)-G(b-s|b)}\right) = E(\dots)\left[\left(\frac{8+N}{11+N}\right)\left(\frac{\alpha_{1,0}+S_1+1}{\alpha_{2,0}+S_2-2}\right) - E(\dots)\right] \qquad (6)$$

where $E(\dots)$ is the first moment from equation 5.

The above approach has two advantages over brute-force approaches to assessment of estimation precision, such as bootstrapping: First, the approach exploits the mathematical structure of the LBV formula, and therefore produces more precise results. In fact, Pham-Gia

(2000) derives the posterior distribution of p_1/p_2, enabling exact posterior inference should anyone need it. Second, the approach is computationally trivial by being just a simple closed-form formula that can be calculated as easily as the estimate itself.

8. Alternative weaker assumption about bidder rationality

Our main LBV formula assumes the bidder prefers the observed bid to every feasible lower bid (equation 1). Consider a weaker assumption inspired by the "zero-intelligence" bidder idea of Gode and Sunder (1993). Suppose the bidder prefers the observed bid to a random lower bid instead of preferring it to every lower bid pointwise. Then, the analogue to equation 1 becomes:

$$G(b)(x-b) > E_{s\in[0,b-R]}[G(b-s)(x-b+s)] \tag{7}$$

and the analogue of equation 3 after still using Assumptions 2 and 3 becomes:

$$x > LBV_0(b) \equiv b + \frac{E_{s\in[0,b-R]}[sG(b-s|b)]}{G(b|b) - E_{s\in[0,b-R]}G(b-s|b)} \tag{8}$$

It is easy to compute this alternative LBV_0, but the simple parametric errors are no longer available.

9. Practical issues in applying the LBV estimator to ad auction data

We now turn to practical issues that arise in an application of our method. There are two issues: first, our LBV computations cannot be applied to bids below reserve because such bids already have the lowest (zero) chance of winning. We cannot somehow ignore such bids since they occur both before and after the switch, and so we conservatively let $LBV(b) = LBV_0(b) = b$ whenever *b<reserve,* and we set the standard error of such LBV to the local average standard error of *LBV*>1. The latter assumption ensures that moving-average standard errors of LBV are undistorted by the conservative augmentation of LBV whenever the bid is below reserve.

The second practical issue is the tradeoff is between "localness" of the competition and precision of the estimate of G because there aren't infinitely many auctions for the same creative occurring exactly at the same time as every focal auction. We solve this tradeoff by requiring at least 100 observations of relevant "nearby" competition for a given focal auction before we compute LBV for the focal bid, and otherwise we again conservatively let $LBV(b) = LBV_0(b) = b$ and the standard error of LBV to the local average. We also make the definition of "nearby" auctions adaptive to the local density of auctions, as we describe next.

The above example in Figure 3 defined auctions "close to the time of the focal auction k" in terms of occurring within two hours. Such a fixed-window definition of proximity identifies thousands of nearby auctions for peak-time auctions of high-volume creatives while not finding enough (at least 100, see previous paragraph) nearby auctions for nighttime auctions of lower-volume creatives. We address this volume-variation issue by allowing the time-window that defines temporal proximity to vary across creatives and time as follows: For any given focal auction, a candidate auction by the same bidder on the same creative is considered *nearby in terms of time* when both of the following conditions are satisfied: 1) it occurred within 6 hours of the focal auction and 2) there occurred fewer than 2000 auctions by the same bidder on the same creative between the time of the focal auction and the time of the candidate auction. This proximity rule automatically expands the time-window during low-volume times while tightening it in busy periods, guaranteeing between 100 and 4000 nearby auctions every time we actually compute the LBV.

After experimenting with different settings of the localness in terms of time, we also decided to expand the definition of localness in terms of bid magnitude to 15 cents. Finally, we implemented the "max" operator in equation 3 only across additional shades that satisfied: 3) additional shade s of at least 2 cents and 4) additional shade s with at least 10 competing bids between b and b-s. Both tuning parameters 3) and 4) ensure that the denominator in the critical ratio (i.e. "p_2" in the notation of the previous section) does not collapse to near zero, which would blow up the standard error. Together, the above "tuning parameters" 1) – 4) deliver estimates of LBV with enough precision to detect insufficient shading given the amount of data we have available.

For each creative, Table 2 in the Appendix documents the proportion of observations for which we computed our *LBV* measures (measured as percentage of pre-switch average valuation), and gives the average sample sizes used in those computations. While there is obvious variation across the creatives, the table shows that on average, our method computes an LBV for about three quarters of the bids, and bases the estimate of G on over 1000 local auctions involving the same creative that occurred within about two and a half hours of the focal auction. For the remaining observations, our method conservatively assigns an *LBV*=100% (i.e. valuations of those impressions are interpreted to be at least the average valuation of an impression before the switch), so any average LBV (for example, we report a 4-hour moving average LBV below) is also conservative.

10. Results

Before summarizing our analysis of all 11 creatives in our data, we return to the three creatives shown in Figure 2, and illustrate the application of our LBV calculation to them in the first subsection. The first subsection will also explain how to interpret our confidence intervals. Then, we turn to a summary of results for all 11 creatives we study.

10.1 Results for three creatives from Figure 2

Figure 4 shows the 4-hour moving average LBV along with conservative *pointwise* 95% confidence intervals. By "pointwise", we mean the local average standard error of a single *LBV* measurement, not the standard error of the moving average. Since the moving average averages over more auctions than just a single LBV computation, it is obviously more precise than the pointwise standard error would suggest. For example, consider the average creative in Table 2 with about 600 auctions an hour and the LBV based on an average of 1000 auctions and a 5-hour window. The 4-hour moving average of such an LBV is based on a 9-hour window containing 5400 auctions, and so its standard error is about $\sqrt{5.4} = 2.3$ times more precise than the pointwise standard error would suggest. As a rule of thumb, one can thus conservatively consider the confidence intervals of the moving average to be at least twice narrower than the pointwise confidence intervals shown in Figure 4.

By comparing the LBV to the pre-switch levels of valuations of the same creative by the same bidder, we can detect insufficient shading whenever the post-switch LBV systematically and significantly exceeds the pre-switch valuation. From this comparison shown in Figure 4, it is clear that B1 did not shade bids down sufficiently until the last month. C2, on the other hand, did not shade bids sufficiently after about the first two weeks. Finally, since the LBV of D1 oscillates around the same magnitude as the pre-switch valuation levels and, we cannot reject the hypothesis that D1 shaded sufficiently.

Figure 4: Lower bounds on valuations the three creatives from Figure 2

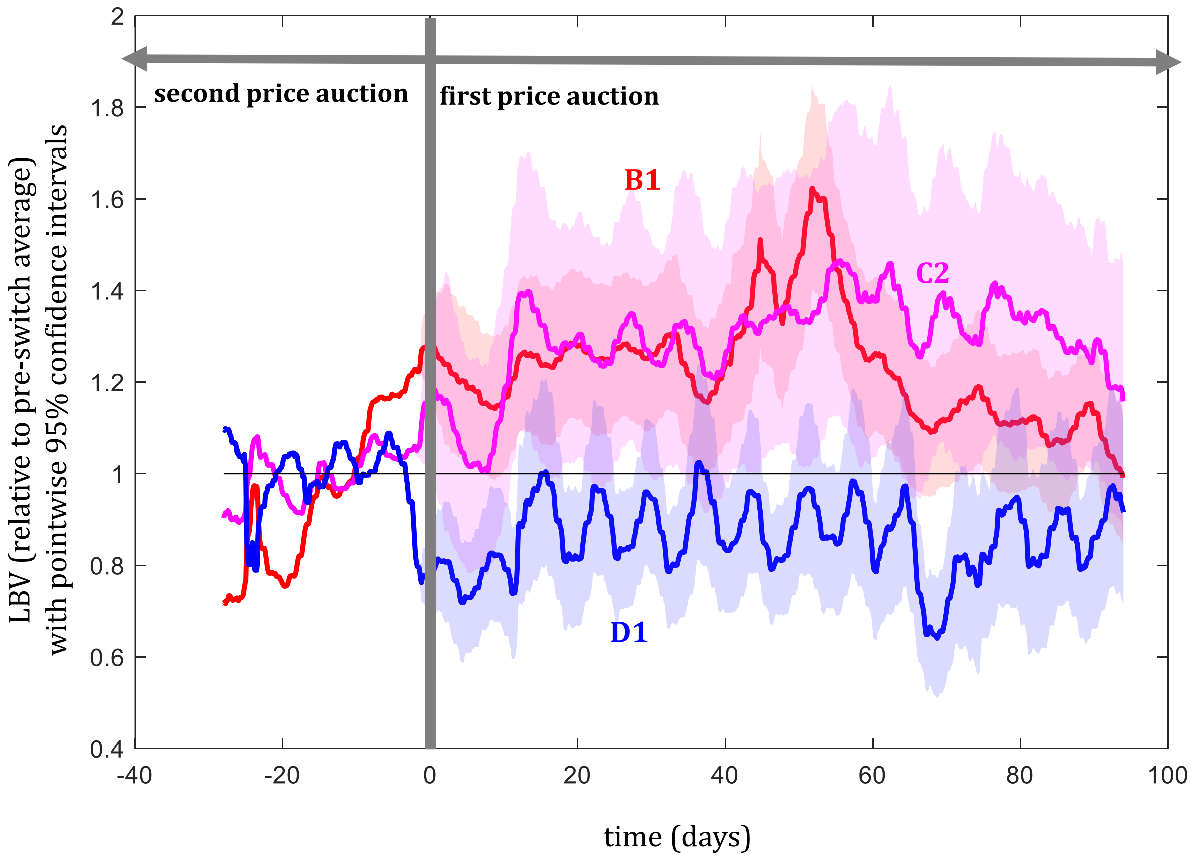

As explained above, the LBV measure isolates insufficient shading given the competition at the time, as opposed to the CPM measure which derives simultaneously from both influences as well as the likely violation of revenue equivalence. To illustrate the difference in practice, consider Figure 5, which shows the two measures along with the bid amount for creative C2. Note that in the tumultuous two weeks after the switch, the LBV and CPM measures are telling different stories: the LBV is consistent with adequate shading while the elevated CPM indicates the competition has increased enough to raise prices compared to pre-switch levels. We conclude that our proposed measure provides information about bidder's bidding strategy above and beyond the prices the bidder pays for impressions.

Figure 5: Cost per impression (CPM) vs. Lower Bound on Valuation for one creative

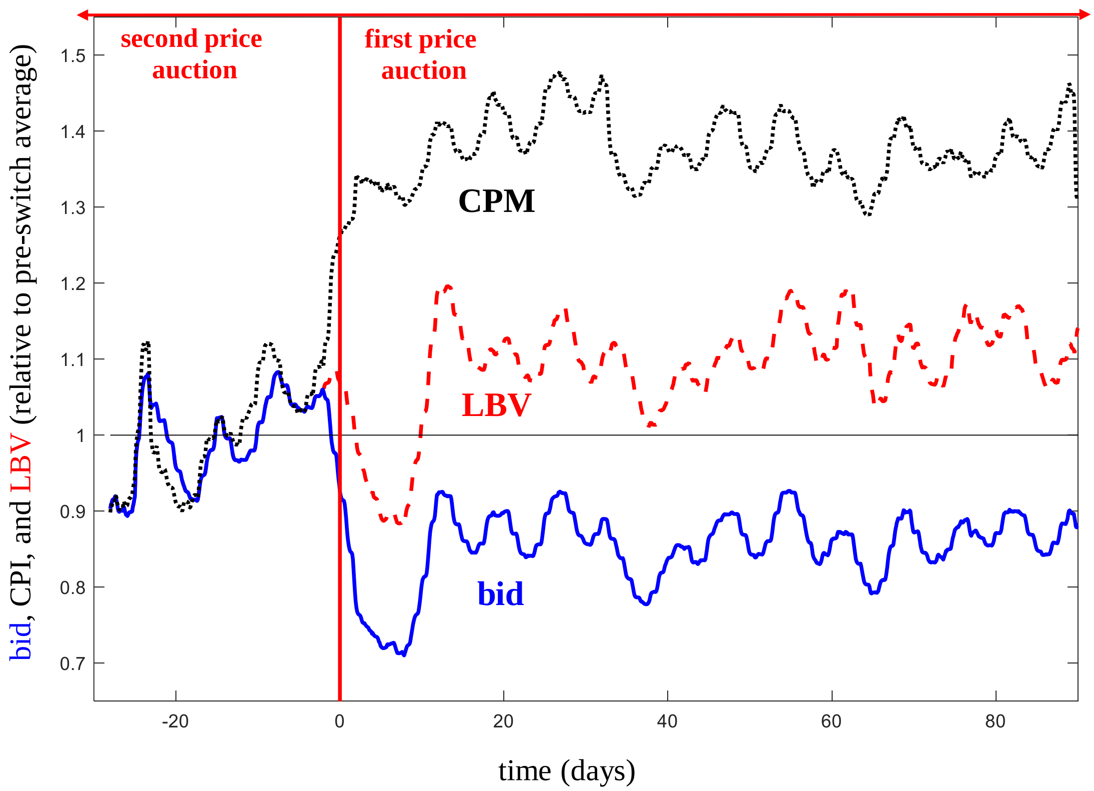

Having described our proposed bound estimator and illustrated its use on three creatives, we now turn to the analysis of the entire dataset.

10.2 Summary of empirical results for all creatives

For each creative in our data, Table 3 shows the average weekly LBV expressed relative to the pre-switch average bid (=valuation) of the same creative. The highlighted cells show creative-weeks in which the average LBV statistically significantly exceeds the average pre-switch valuation at the 5% level given the average pointwise standard error that week. This statistical inference measures the chance that the LBV in one randomly-selected auction that has the average LBV for that week significantly exceeds the pre-switch valuation given the average magnitude of single-auction standard errors that week. For the same reason as explained in the previous section for the case of a 4-hour moving average, this pointwise inference is conservative in that the standard error of the average values shown in Table 3 is clearly much smaller than the average pointwise error.

Table 3: Relative LBV by week and creative

week	A1	B1	B2	B3	B4	C1	C2	D1	D2	D3	D4	average
-3	91%	83%	85%	84%	96%	90%	96%	101%	90%	96%	98%	92%
-2	84%	102%	99%	111%	108%	103%	100%	103%	86%	112%	101%	101%
-1	105%	120%	116%	111%	109%	109%	105%	99%	122%	94%	101%	108%
0	122%	123%	132%	116%	152%	109%	112%	77%	162%	87%	121%	119%
1	233%	120%	135%	120%	160%	126%	121%	86%	162%	55%	119%	131%
2	167%	125%	151%	121%	116%	140%	129%	88%	133%	63%	115%	123%
3	123%	127%	153%	121%	122%	136%	128%	89%	154%	82%	116%	123%
4	118%	127%	131%	122%	127%	135%	128%	87%	144%	127%	134%	125%
5	155%	119%	121%	115%	152%	118%	126%	93%	129%	91%	125%	122%
6	156%	139%	197%	110%	145%	126%	134%	87%	123%	97%	160%	134%
7	144%	154%	164%	108%	142%	132%	138%	91%	131%	148%	187%	140%
8	157%	127%	114%	99%	207%	145%	141%	90%	117%	174%	191%	142%
9	163%	112%	104%	85%	206%	140%	132%	80%	115%	147%	191%	134%
10	182%	116%	109%	87%	218%	139%	134%	82%	131%	139%	182%	138%
11	157%	111%	103%	83%	197%	142%	136%	87%	131%	136%	192%	134%
12	152%	109%	132%	82%	192%	135%	128%	88%	128%	138%	182%	133%

It is immediately evident that the D1 creative is in a minority, joined only by B2 and B3 towards the end of the sample period as an example of a creative, for which our method does not detect insufficient shading. Interestingly, the general trend is not toward sufficient shading: creatives B4 and D3 start out with seemingly adequate shading during the first month after the switch, but fail to shade adequately towards the end of the data period.

Looking back at Table 1, B4 is also another example of a creative with CPM telling a different story from LBV: while the CPM jumps to above 200% after the switch and stays there, the bid-shading does not seem significantly inadequate during weeks 2-4, and then becomes so towards the end of the data period. Summarizing across all 11 creatives, the average LBV remains about 30 percent above the pre-switch average valuation throughout the post-switch period.

Figure 6: Lower Bound on Valuations over time, average across creatives

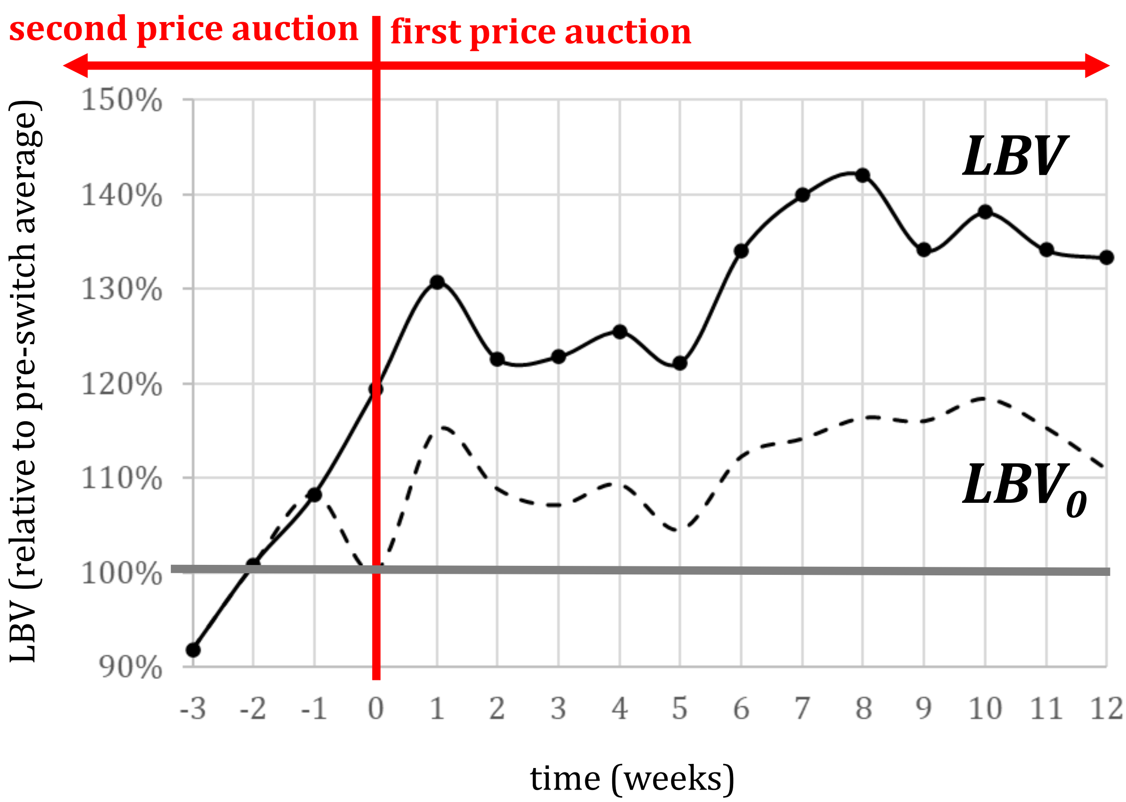

In addition to suggesting widespread insufficient bid shading even months after the switch, there is no detectable downward trend in the average LBV over time. Figure 6 plots the average LBV by week shown in the last column of Table 3 along with the analogue for LBV_0 reported in detail in Table 4 in the Appendix. The latter metric indicates that even if we assume that each bidder preferred its bid to a random lower bid (as opposed to a "cherry-picked" lower bid that maximizes our bound), we still find evidence of insufficient shading throughout the three months after the switch. For both metrics, the average lines in Figure 5 represent a wide range of behaviors of the different creatives shown in Table 3.

In summary, we conclude that the bidders we study took more than three months to adjust to the new pricing rule if they ever adjusted at all. Thanks to the insufficient bid-shading, advertisers paid higher prices for impressions during all those months. At least part of the average 35 percent increase in CPM evident from Table 1 can thus be attributed to insufficiently shaded bidding strategies of the bidders associated with the creatives, and not merely to increased competition due to new entry or insufficient shading by other bidders participating in the auctions we study or to a failure of revenue equivalence.

Every one of the three multi-creative bidders in our data has at least one creative with consistently insufficient shading throughout the data period, so we do not find evidence of heterogeneity in bidding sophistication across bidders.

11. Alternative explanations

We can rule out several alternative explanations of the data patterns we document: wrong pre-switch valuation definition, trend seasonality, drift seasonality, and forward-looking strategic behavior. In this section, we briefly discuss our evidence against these explanations one at a time.

11.1. Wrong pre-switch valuation definition

Several of the creatives in Table 3 exhibit a large variance in pre-switch average weekly bids. If we conservatively take the highest pre-switch average weekly bid as the estimate of pre-switch valuation of the impression, most of the LBV estimates are no longer pointwise significantly above new pre-switch valuations. However, the (quite precisely estimated) weekly average LBVs exceed the new pre-switch valuations by at least 10 percent in five of the creatives (A1, B4, C1, C2, and D4) consistently throughout the entire post-switch period, and two others (B2 and D3) for more than half of the post-switch weeks. This conservative definition of the pre-switch average thus adds only two creatives (B2 and D1) to the list of creatives for which we do not detect insufficient shading for most of the post-switch period.

11.2. Trend seasonality

Advertising expenditures were gradually rising during the time of the switch, and some industry observers even suggest that there is a predictable yearly cycle whereby advertising spending builds from winter to spring and summer. It is therefore possible that our evidence of bidders bidding as if their valuations were higher was due to valuations actually rising over time. One test of this explanation is already available by considering the pattern of detected insufficient shading in Table 3: a gradual rise would make more and more creatives seem insufficiently shaded over time. Instead, the number of such creatives is 7 or 8 in most weeks without any discernible trend. The average LBV in the rightmost column of Table 3 and solid line in Figure 6 also does not exhibit any upward trend.

To further test the trend seasonality explanation of the seemingly higher valuations after the switch, we regressed the relative bid (measured in percent of average valuation before switch such that the number 100 denotes 100%) on time (measured in days) starting one week after the switch for each creative separately. We chose to include only bids after the switch so as not to confound the gradual time trend with the effect of the switch, and chose to start one week after the switch to allow for some adjustment and skip the turbulent first week. The remaining 3 months of data should be plenty to detect a systematic upward trend. Table 4 shows the slopes of relative bids in time. There is no systematic upward trend in that roughly half (5) of the slopes are positive and the remainder (6) negative. The average slope is +0.028% per day – hardly enough to increase valuations by over 30 percent even at the end of the data period (less than 100 days later). Finally, the more positive slopes do not necessarily correspond to more detection of insufficient shading: A1, B1 and D2 all have negative slopes and insufficient shading. C2 has the closes slope to zero, yet it exhibits an LBV consistently above unity throughout.

Table 4: Linear regressions of post-switch bids starting with week 1 on time

Creative	relative bid slope per day ("1"=1% bid before)	95% confidence interval		R^2	Num obs
A1	-0.048	-0.056	-0.040	0.000	702,495
B1	-0.098	-0.104	-0.092	0.002	627,916
B2	-0.326	-0.328	-0.323	0.013	5,251,840
B3	-0.349	-0.356	-0.341	0.020	408,187
B4	0.231	0.224	0.238	0.036	107,996
C1	0.027	0.024	0.031	0.000	2,854,782
C2	0.002	-0.002	0.005	0.000	2,407,707
D1	-0.025	-0.030	-0.020	0.000	190,415
D2	-0.356	-0.365	-0.347	0.048	117,891
D3	0.757	0.754	0.760	0.159	1,209,302
D4	0.497	0.495	0.499	0.528	280,194

11.3. Drift seasonality

It is possible that valuations drift over time, following something like a random walk with a random trend – some upward and some downward – and we are erroneously ascribing the creatives than happened to drift upward with insufficient bid shading. One piece of evidence against this alternative explanation is analogous to the above argument against trend seasonality: the number of creatives with insufficient shading does not seems to increase over time. A more direct test of creatives drifting apart is in Figure 7 which plots the standard deviation across creatives in the moving-average bid (relative to pre-switch average) in two-hour intervals. If the creatives were drifting apart, this standard deviation should be increasing over time. Instead, Figure 7 shows a period of initial volatility after the switch lasting about two weeks, followed by almost three months of much lower volatility without any discernable trend.

Figure 7: standard deviation across creatives in the moving average relative bid

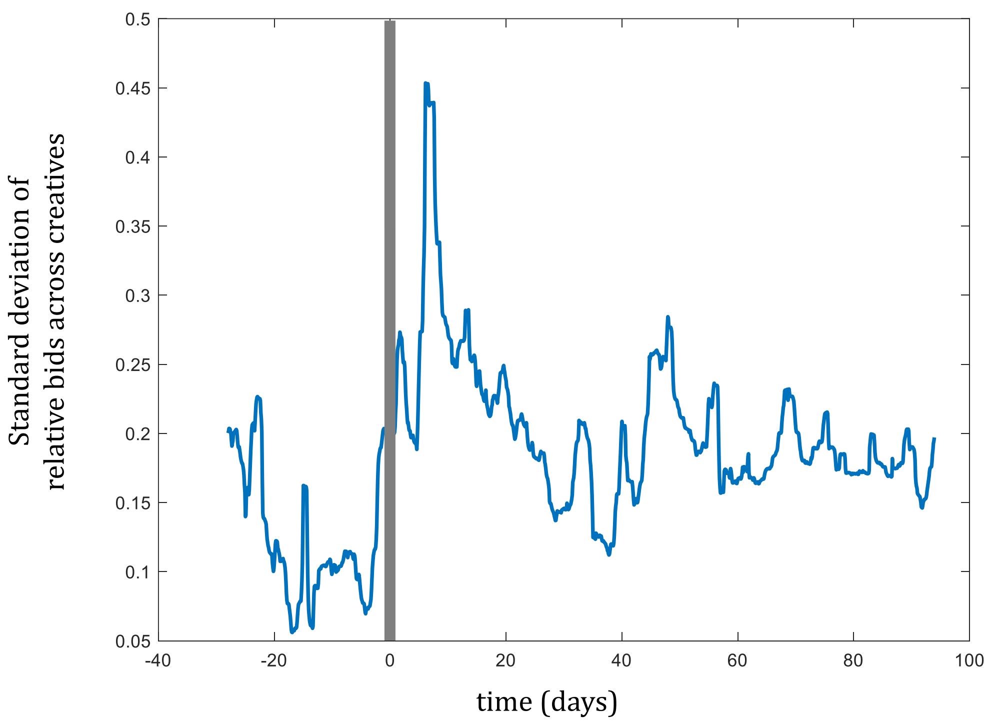

11.4. Forward-looking increase in net valuation given higher CPM levels

When it becomes more expensive to show the same ad in the near future, as it did after the switch to first-price rules, then the net value of showing the ad now rises even if the profit to the advertiser from showing the ad remains the same. The net valuation relevant for bidding increases because the opportunity cost of winning (equal to surplus from participating in near-future auctions) falls. Did the bidders we study realize the higher cost of advertising and correctly increased their net valuations of showing the same ad? We cannot see what the bidders were thinking, but we can test for other behaviors consistent with such hyper-rational behavior. One obvious prediction is that bidders who realize the cost of advertising increased buy less of it. Fixed budgets deliver this prediction mechanically, and rational behavior delivers it because demand tends to be downward sloping. In contrast with this prediction, we do not see a reduction in advertising volume after the switch. Instead, the purchase volume of each creative we study rises for at least two months after the switch. Table 5 shows that most of the creatives initially participated in more auctions than before the switch, and only brought down their entry after about two months of winning about twice more impressions than before the switch.

The above entry and winning patterns are inconsistent with a strategic reduction in ad purchasing due to higher prices. Instead, the pattern suggests that competition weakened after the switch, either due to more shading or due to reduced participation.[7] Sensing an opportunity, the algorithms behind the long-running creatives we study rushed in to capitalize on the opportunity, but generally ended up buying too much advertising due to insufficient bid-shading. Note, however, that merely buying more ads is not necessarily evidence of insufficient shading: creative D1 exhibits the largest volume increase despite its apparently sufficient shading. In general, Table 5 shows a lot of variation in volume changes, both across creatives and over time. This variation is doubtless caused by adjustments in

[7] Reduced participation due to shifting spending to other publishers is the more plausible explanation. A recent blog post explains the mechanism eloquently: "DSPs primary job is to better the ROI, not the other way around. Hence, as soon as, they realize bidding on your inventory isn't helping the buyer to achieve his/her goal (ROI, CPA, etc.), they'll move the spending to a different publisher."(Automatad 2021).

budgets and other campaign settings idiosyncratic to particular creatives at particular times for reasons unobserved to us.

Table 5: Participation and purchase volume after switch, by month

Creative	# auctions entered (% of month before switch)			# auctions won (% of month before switch)			winning probability (% of month before switch)		
	1 after	2 after	3 after	1 after	2 after	3 after	1 after	2 after	3 after
A1	97%	148%	45%	84%	142%	62%	86%	96%	136%
B1	233%	228%	105%	256%	368%	140%	110%	161%	134%
B2	212%	210%	105%	299%	373%	130%	141%	177%	124%
B3	225%	237%	109%	269%	333%	114%	120%	140%	104%
B4	377%	311%	236%	307%	376%	396%	81%	121%	168%
C1	112%	111%	56%	100%	115%	67%	89%	104%	119%
C2	119%	135%	65%	94%	130%	66%	79%	96%	100%
D1	387%	503%	159%	369%	571%	280%	95%	113%	176%
D2	148%	126%	66%	233%	213%	143%	158%	169%	216%
D3	203%	143%	72%	140%	161%	113%	69%	113%	158%
D4	102%	26%	40%	173%	101%	176%	169%	390%	443%
average	**201%**	**198%**	**96%**	**211%**	**262%**	**153%**	**109%**	**153%**	**171%**
median	**203%**	**148%**	**72%**	**233%**	**213%**	**130%**	**95%**	**121%**	**136%**

Table 5 also shows the evolution of winning probability (i.e. the ratio of winning to participation), which gradually rose from the pre-switch level for most creatives throughout the post-switch period. This pattern is consistent with insufficient shading in the long run, and explains why even after adjusting participation frequency far below pre-switch levels, most creatives ended up with more volume (than before switch) even three months after the switch.

In summary, our explanation that the bidders are not paying enough attention to their bidding strategy and not shading sufficiently is consistent with the patterns shown in Table 5, but the alternative explanation of seemingly higher valuations due to hyper-rational forward-looking expectation of rising advertising costs is not.

12. Discussion

We analyzed the response of bidders (demand-side platforms) to a switch in auction pricing rules from second-price to first-price on one online advertising auction platform in 2019. In terms of overall publisher revenue and ad prices, we find a transitory increase lasting four weeks. But when we zoom in on the behavior of concrete bidders bidding to show long-running creatives, we find evidence that bidders generally struggled to adjust to the switch for at least three months. Specifically, the bidders bid on the creatives we analyzed *as if* the switch in pricing rules suddenly increased their valuations of an impression by at least 30 percent for at least three months after the switch. In other words, the bidders we observe actively bidding to show the same creative in the same space throughout the study period did not shade their bids enough in response to the switch. Therefore, the transitory nature of the revenue and price increase is not necessarily due to an adjustment of bidding strategies, at least not among the bidders we study. It seems other adjustments must have also taken place, for example an adjustment of bidding strategies on only new creatives by the bidders we study, an excessive adjustment of bidding by some other bidders, or a reduction in entry suggested by Automatad (2021).

We can reach the above substantive “as if” conclusion thanks to a new bound estimator we developed specifically for the purpose of measuring insufficient bid shading. Building on classic work in the econometrics of first-price auctions, our approach derives a lower bound on valuation in each post-switch auction from bid magnitude and the intensity of local competition. By considering local competition not only in the sense of time but also in the sense of bid magnitude, our estimator accounts for affiliation in valuations of impressions – a natural assumption in the online advertising context. By considering the actual competition separately for every auction and bidder, our estimator does not require bidder symmetry or equilibrium bidding to hold.

Armed with an estimate of the lower bound on valuation in each post-switch auction, we then compare the distribution of the bounds with the distribution of pre-switch valuations revealed directly as bids (thanks to second-price rules). When the lower bound on post-switch valuations exceeds the pre-switch valuations of showing the same creative in the same location on the publisher's website, we conclude that the bidder did not shade the first-price bid sufficiently. We reach this conclusion in 8 of the 11 creatives we study. Every one of the three multi-creative bidders in our data has at least one creative with consistently insufficient shading throughout the three months after the switch covered by our data, so we do not find evidence of heterogeneity in bidding sophistication across bidders.

The insufficient shading has profound implications on the cost of advertising to advertisers we study. For the average long-running creative, costs per impression (CPM) rose by about 35 percent after the switch and remained elevated for months. Note that this does not immediately imply that the revenue of the bidding platform remained elevated as well - our sample of creatives is selected to address our question of interest, not to be representative of the market as a whole, and we do not measure entry of other bidders and creatives. However, the CPM increase does mean that our bidders had to pay more to show these long-running creatives after the switch. The bound estimator discussed above allows us to attribute at least part of this price increase to insufficient shading by frequently participating bidders as opposed to an increase in competition after the switch or a violation of revenue equivalence. Note also that even if revenue equivalence holds in equilibrium, it may not hold for a market in transition such as the one we study: similar to the argument in

Deltas and Engelbrecht-Wiggans (2001), it may be that irrationality of the competing bidders increases costs even for a bidder who shades bids correctly.

When we look at the lower bound on valuations over time, we do not find a detectable downward trend in the three months after the switch. We conclude that the bidders took more than three months to adjust to the new pricing rule if they ever adjusted at all. The bid-adjustment process, while heterogeneous, was thus surprisingly slow given the apparent technological capabilities the industry.

Since the switch we study occurred early in the calendar year, it is possible that the elevated LBV we find is a result of a seasonal increase in impression valuations over time. However, such an increase would be gradual, not abrupt right after the switch. As part of our exploration of alternative explanations, we did not find any evidence of a gradual upward trend in bids. Another alternative inconsistent with our data is that the net valuations of the same impression actually increased after the switch due to forward-looking bidders correctly anticipating higher costs. Such hyper-rational bidders would also buy less advertising given its higher cost, but the bidders we study actually bought more exposures immediately after the switch.

One way to interpret the lack of adjustment we document is that first-price auctions were still relatively new to the DSPs at the time of the switch we study (early 2019), so coming up with optimal bidding strategies required more effort than it does today. Anecdotal evidence from the time period suggests that most of the bidders we study did not even actively take the reserve price into account when formulating their bids – something all DSPs do routinely today. Taking the reserve price into account is a clear acknowledgement of first-price rules because it is necessary for bid optimization in first-price auctions, but not in the legacy second-price auctions that have dominant strategies. We have no doubt that bidding strategies have gotten a lot more sophisticated in the few years since our sample, so our paper should be viewed as a study of adjustment to a pricing mechanism with which bidders are not too familiar, not as a critique of today's industry participants. Goke et al (2022) indeed find some evidence that ad-price adjustment to first-price rules has been quicker in auctions that switched more recently.

The insufficient shading we document means that valuation estimates using standard econometrics of first-price auctions (e.g. Athey and Haile 2002) on current data from real-time bidding (RTB) markets for display advertising may be biased upwards. Unlike in the case of human bidders, where risk-aversion seems to fix this bias (Bajari and Hortaçsu 2005), it is not clear why RTB algorithms exhibit the upward bias in inferred valuations. More research is needed to figure out whether the adjustment we looked for eventually happened, and what was the underlying cause of the insufficient adjustment.

Assuming that the strategy adjustment eventually happened, its slow speed we document implies that analysts of real-time bidding on online display advertising should not rely on short-run A/B tests when evaluating the profitability of different auction pricing rules. Davies (2019) describes a prominent example of such a test when she says "Google has spent the last few months testing the outcome of running first-price auctions across 10% of its Google Ad Manager inventory." Our results imply that such a test of the first-price vs the second-price auction rule is likely to wildly overestimate the long-run profitability of the first-price rule because it does not allow sufficient time for the surprisingly slow adjustment in bidding strategies to occur. More work on bidder learning and other adjustment is needed to correctly extrapolate short-run tests into long-run predictions.

Appendix: additional figures, tables, and explanations

Figure A1: Naïve vs. correct counter-factual probability of winning

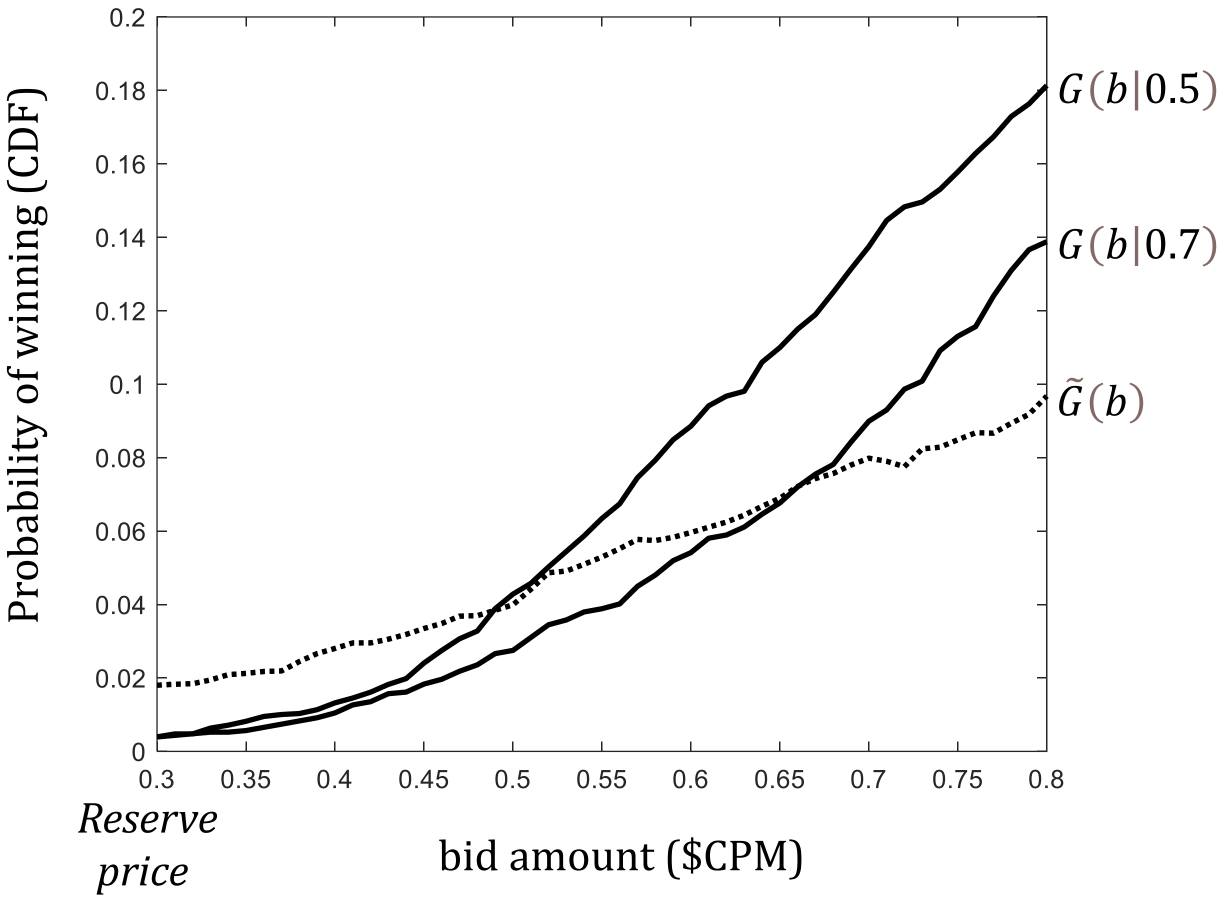

Note to Figure: The dashed line shows the empirical probability of winning at different bid levels for one creative during one day two weeks after the switch, and it is based on over ten thousand observations. The two solid lines show the correct counter-factual probability of winning when the bidder's own bid is 0.5 (higher line) and 0.7 (lower line). The correct counter-factuals are the cdf of competing bids in the subset of auctions on the same day in which the bidder's own bid was within 10 cents of the focal bid.

Table 2: Implementation details of the LBV computation

		Classification of observed bids according to practical issues in LBV computation			Properties of observations used in computing LBV > 1	
Creative	Number of bids	bids above reserve with at least 100 nearby auctions	bids above reserve with fewer than 100 nearby auctions	bids below reserve	average number of nearby auctions	average size of "nearby" window (hours)
		→ LBV > 1	→ LBV = 1	→ LBV = 1		
A1	968,238	68%	2%	30%	1283.9	3.9
B1	778,412	74%	3%	23%	1132.8	6.5
B2	6,612,124	69%	2%	29%	1447.2	1.5
B3	505,714	64%	5%	31%	1065.4	7.4
B4	123,187	73%	27%	1%	452.9	6.7
C1	4,253,188	53%	2%	45%	871.4	2.6
C2	3,254,961	61%	2%	37%	865.9	3.1
D1	214,450	83%	4%	13%	763.3	7.7
D2	159,370	79%	20%	1%	374.8	7.5
D3	1,612,881	76%	2%	22%	1317.3	4.3
D4	487,630	99%	0%	1%	2259.0	7.5
Average	**1,724,560**	**73%**	**6%**	**21%**	**1075.8**	**5.3**

Example of how using a valuation net of future opportunities reduces the dynamic problem to a static one

Suppose the bidder in fact faces a rapid (no time discounting) infinite sequence of opportunities, and only wants to obtain a single impression ("unit demand" - a simple and tractable example of a budget constraint). Such a bidder solves the following dynamic program for the NPV of having a unit-demand valuation of v: $V(v) = \max_b G(b)(v - b) + [1 - G(b)]V(v)$.

It is clear that V solves: $0 = \max_b G(b)[v - V(v) - b]$, so the bidder bids as if he were in a single auction and his valuation were net of the option value of losing: $x \equiv v - V(v)$. See Milgrom and Weber (2000) for additional analysis of this issue.

Table 4: LBV$_0$: lower bound on valuations under weaker assumption about bidder rationality, by week and creative

week	A1	B1	B2	B3	B4	C1	C2	D1	D2	D3	D4	median
0	66%	114%	122%	112%	109%	57%	87%	93%	146%	75%	120%	109%
1	163%	114%	123%	115%	130%	101%	91%	93%	153%	62%	122%	115%
2	132%	115%	129%	117%	99%	106%	93%	87%	125%	68%	126%	115%
3	115%	115%	129%	115%	95%	104%	92%	88%	129%	75%	121%	115%
4	114%	115%	119%	114%	105%	102%	91%	84%	128%	103%	128%	114%
5	130%	112%	115%	111%	117%	58%	91%	91%	110%	82%	132%	111%
6	129%	119%	146%	106%	114%	98%	92%	91%	103%	82%	156%	106%
7	123%	132%	126%	106%	115%	98%	90%	91%	109%	104%	163%	109%
8	133%	113%	114%	104%	121%	103%	93%	95%	104%	128%	173%	113%
9	135%	109%	111%	94%	155%	103%	91%	91%	106%	111%	171%	109%
10	145%	110%	112%	99%	165%	103%	91%	91%	117%	102%	166%	110%
11	130%	110%	109%	76%	153%	108%	95%	103%	115%	102%	168%	109%
12	125%	106%	115%	67%	138%	103%	90%	99%	111%	104%	163%	106%

Note to Table: Standard errors are not readily available. The red shaded cells indicate the weeks when a given creative's average LBV$_0$ exceeds 110% of the pre-switch valuation of the same creative.

References

Athey Susan, Bayati M., Doudchenko N., Imbens G., Khosravi K. (2017) "Matrix completion methods for causal panel data models". arXiv preprint arXiv:1710.10251.

Athey, Susan and Philip Haile (2002) Identification of Standard Auction Models. Econometrica 70(6): 2107-2140.

Automatad Team (2021). Bid Shading – Who Offers it and How it Impacts Publishers? https://headerbidding.co/bid-shading/.

Bajari, Patrick and Ali Hortaçsu (2005) Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data. *Journal of Political Economy* 113 (4): 703-41.

Battigalli, Pierpaolo, and Marciano Siniscalchi. "Rationalizable bidding in first-price auctions." Games and Economic Behavior 45.1 (2003): 38-72.

Bekker-Nielsen Dunbar M, Finnie TJR, Sloane B, Hall IM (2019) Methods for calculating credible intervals for ratios of beta distributions with application to relative risks of death during the second plague pandemic. *PLoS ONE* 14(2): e0211633.

Doraszelski, Ulrich, Gregory Lewis, and Ariel Pakes (2018) Just Starting Out: Learning and Equilibrium in a New Market. *American Economic Review* 108 (3), 565–615.

Gobillon L. and Magnac T. (2016) "Regional Policy Evaluation: Interactive Fixed Effects and Synthetic Controls".*The Review of Economics and Statistics*, July 2016, Vol. 98, No. 3, pp. 535–551.

Hortaçsu, Ali and David McAdams (2010) Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market *Journal of Political Economy* 118(5), 833-865

Choi, Hana and Carl. F. Mela (2019). Display advertising pricing in exchange markets. Working paper.

Choi, Hana and Carl F. Mela and Santiago R. Balseiro and Adam Leary (2020) Online Display Advertising Markets: A Literature Review and Future Directions. *Information Systems Research* 31 (2): 556-575.

Cox, James C., Vernon L. Smith, and James M. Walker (1988) Theory and Individual Behavior of First-Price Auctions. *Journal of Risk and Uncertainty* 1 (1), 61-99.

Davies, Jessica (2019). "What to know about Google's implementation of first-price ad auctions", posted on Digiday.com on Septermber 6, 2019. https://digiday.com/media/buyers-welcome-auction-standardization-as-google-finally-goes-all-in-on-first-price/

Deltas, George & Engelbrecht-Wiggans, Richard (2001) Auctions with an inexpert bidder. *Economics Letters* 73(1): 35-42.

Despotakis, Stylianos and R Ravi and Amin Sayedi (2021) First-price auctions in online display advertising. *Journal of Marketing Research*. Forthcoming.

Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero intelligence traders:
Market as a partial substitute for individual rationality. *Journal of Political Economy* 101: 119–137.

Goke, Shumpei, Gabriel Y. Weintraub, Ralph Mastromonaco, and Sam Seljan (2022) Bidders' Responses to Auction Format Change in Internet Display Advertising Auctions. Working paper available at SSRN 4021221 (2022).

Guerre, Emmanuel and Isabelle Perrigne and Quang Vuong (2000). Optimal Nonparametric Estimation of First-Price Auctions. *Econometrica* 68(3): 525-574.

Haile, Philip and Elie Tamer (2003) Inference with an Incomplete Model of English Auctions *Journal of Political Economy* 2003 111(1): 1-51.

Li, Tong, Isabelle Perrigne, and Quang Vuong (2002) Structural Estimation of the Affiliated Private Value Auction Model. *The RAND Journal of Economics* 33(2): 171-93.

Milgrom, Paul R. and Robert Weber (2000). A theory of auctions and competitive bidding, II P.D. Klemperer (Ed.), The Economic Theory of Auctions, vol. 1, Edward Edgar, Cambridge, UK.

Nekipelov, Denis and Vasilis Syrgkanis, and Eva Tardos (2015) Econometrics for Learning Agents. Proceedings of the Sixteenth ACM Conference on Economics and Computation. June 2015, Pages 1–18.

Pham-Gia, T. (2000) Distribution of the Ratios of Beta Variables. *Communications in Statistics—Theory and Methods* 29 : 2693-2715.
Rafieian, Omid (2019). Revenue-optimal dynamic auctions for adaptive ad sequencing. Working paper.

Sluis, Sarah (2019). Everything You Need To Know About Bid Shading. *AdExchanger.* (https://www.adexchanger.com/online-advertising/everything-you-need-to-know-about-bid-shading/) .

Xu, Y. (2017). "Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models". Political Analysis, 25(1), 57-76. doi:10.1017/pan.2016.2

Xu, Y. and Licheng L. (2022). "gsynth: Generalized Synthetic Control Method". R package version 1.2.1

Zeithammer, Robert (2019) Soft Floors in Auctions. *Management Science* 65(9): 3949-4450